arXiv:1605.04877v1 [math.CO] 16 May 2016

BOREL VERSION OF THE LOCAL LEMMA

ENDRE CSÓKA, ŁUKASZ GRABOWSKI, ANDRÁS MÁTHÉ, OLEG PIKHURKO, KONSTANTINOS TYROS

Abstract. We prove a Borel version of the local lemma, i.e. we show that, under suitable assumptions, if the set of variables in the local lemma has a structure of a Borel space, then there exists a satisfying assignment which is a Borel function. The main tool which we develop for the proof, which is of independent interest, is a parallel version of the Moser-Tardos algorithm which uses the same random bits to resample clauses that are far enough in the dependency graph.

Contents

1. Introduction 2
1.a. Basic notation and definitions 6
2. Moser-Tardos algorithm 7
2.a. Borel graphs 7
2.b. Moser-Tardos algorithm with limited randomness 8
3. Preliminaries on landscapes of trees 11
3.a. Landscapes 11
3.b. Grounded landscapes 12
3.c. Counting grounded landscapes 13
3.d. The b-valued sequences encoded by a decorated landscape 16
3.e. Landscape restrictions 16
4. Local Lemma on Borel graphs of subexponential growth 17
4.a. Analysis of the Moser-Tardos algorithm 17
4.b. Borel Local Lemma 21
Appendix A. Some standard lemmas 23
References 26

1. Introduction

Conventions. Throughout the article we frequently use the fact that any natural number is equal to the set of all smaller natural numbers. In particular a natural number is not only an element of $\mathbb{N} := \{0, 1, \ldots\}$ but also a subset of $\mathbb{N}$.

A vertex in a graph is allowed to have at most one self-loop. The *neighbourhood* of a vertex is the set of all vertices connected to that vertex by an edge. The *degree* of a vertex is the cardinality of its neighbourhood. In particular the degree of a vertex with a self-loop and no other edges is equal to 1.

Local Lemma. Let us start by recalling a version of the local lemma. The first version of the local lemma was proved by Erdős and Lovász [3]. The version we present follows from the subsequent improvement of Lovász (published by Spencer [9]). For more historical context we refer to the classical exposition in [1].

Let G be a graph and let b be a natural number greater than 1. The elements of the vertex set $V(G)$ should be thought of as variables which can take values in the set $b = \{0, 1, 2, \ldots, b-1\}$. Let $\mathbf{R}$ be a function whose domain is $V(G)$ and such that for $x \in V(G)$ we have that $\mathbf{R}(x)$ is a set of b-valued functions defined on the neighbourhood of x. Such a function $\mathbf{R}$ is an example of a *local rule* on G. We say that $f\colon V(G) \to b$ *satisfies* $\mathbf{R}$ if for every $x \in V(G)$ the restriction of f to the neighbourhood of x belongs to $\mathbf{R}(x)$.

The local lemma gives a condition under which a satisfying assignment exists. For $x \in V(G)$ let $p(x)$ be the probability of failure at x, i.e. $p(x) := 1 - \frac{|\mathbf{R}(x)|}{b^{\deg(x)}}$, where $\deg(x)$ is the degree of x. Let $\mathrm{Rel}(G)$ be the graph whose vertex set is $V(G)$ and such that there is an edge between x and y if and only if the neighbourhoods of x and y are not disjoint (we allow x and y to be equal). Finally, let Δ be the maximal vertex degree in $\mathrm{Rel}(G)$.

Theorem 1.1 (Lovász's Local Lemma [9])**.** *If for all $x \in V(G)$ we have $p(x) < \frac{1}{e\Delta}$ then there exists $f\colon V(G) \to b$ which satisfies $\mathbf{R}$.*

In order to motivate our *Borel version* of Theorem 1.1, let us recall a classical application of the local lemma to colorings of Euclidean spaces from [3]. A *b-coloring* of $\mathbb{R}^n$ is a function $f\colon \mathbb{R}^n \to b$. We say that a set $U \subset \mathbb{R}^n$ is *multicolored* with respect to a b-coloring if U contains points of all b colors.

Corollary 1.2 ([3])**.** *Let $T \subset \mathbb{R}^n$ be a finite set of vectors and let $b \in \mathbb{N}$ be such that $b(1 - \frac{1}{b})^{|T|} < \frac{1}{e \cdot |T|^2}$. Then there exists a b-coloring of $\mathbb{R}^n$ such that for every $x \in \mathbb{R}^n$ the set $x + T$ is multicolored.*

Proof. We define a graph G as follows. Let $V(G) := \mathbb{R}^n$ and let $(x, y) \in E(G)$ if for some $t \in T$ we have $x + t = y$. We define a local rule $\mathbf{R}$ by letting $\mathbf{R}(x)$ be the set of all surjections from the neighbourhood of x to b. It is clear that in order to prove the corollary we need to find a function $f\colon V(G) \to b$ which satisfies $\mathbf{R}$.

It is easy to bound the maximal degree in $\mathrm{Rel}(\mathcal{G})$ by $|T|^2$. Since the image of a non-surjection is contained in a set with $b-1$ elements, we see that for all $x \in V(G)$ the probability $p(x)$ of failure is bounded from above by $\frac{b(b-1)^{|T|}}{|b^T|} = b(1-\frac{1}{b})^{|T|}$. Thus, we can apply Theorem 1.1 to finish the proof. □

Borel Local Lemma. Our Borel version of the local lemma allows to deduce that the b-coloring in Corollary 1.2 can be demanded to be a Borel function.

Let G be a graph as before, but let us additionally assume that $V(G)$ is a standard Borel space and that $\mathbf{R}$ is a *Borel local rule*. Since it is notationally involved to precisely define what it means for a local rule to be Borel, we defer doing it to Section 2. For now we only assure the reader that in all naturally occurring applications of the local lemma the local rules are in fact Borel. For example, the local rule in the proof of Corollary 1.2 is Borel.

We say that a graph G is of *uniformly subexponential growth* if for every $\varepsilon > 0$ there exists r such that for all $R > r$ and all vertices $v \in V(G)$ the set of vertices of G at distance at most R from v has cardinality less than $(1+\varepsilon)^R$.

Theorem 1.3 (Borel Local Lemma)**.** *Let G be graph such that $V(G)$ is a standard Borel space and let $\mathbf{R}$ be a Borel local rule on G. Furthermore let us assume that the graph $\mathrm{Rel}(G)$ is of uniformly subexponential growth, and let Δ be the maximal degree in $\mathrm{Rel}(G)$.*

If for all $x \in V(G)$ we have $p(x) < \frac{1}{e\Delta}$ then there exists a Borel function $f\colon V(G) \to b$ which satisfies $\mathbf{R}$.

This theorem follows from a more general Corollary 4.8.

In the proof of Corollary 1.2 the graph $\mathrm{Rel}(G)$ is easily seen to be of uniformly subexponential growth. Thus we obtain the following corollary of Theorem 1.3.

Corollary 1.4. *In Corollary 1.2 we can additionally demand the b-coloring of $\mathbb{R}^n$ to be a Borel function.*

Remarks 1.5. (i) The assumption of the local lemma for example in [9, Theorem 1.5] is that $p(x)$ should be less that $\frac{1}{e(\Delta+1)}$, as opposed to our $\frac{1}{e\Delta}$. The difference is only notational, and comes from the fact that we count self-loops when calculating vertex degrees.

(ii) One can weaken the assumption on $p(x)$ in the local lemma to $p(x) < \frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}$, and this is best possible [7]. Corollary 4.8 shows that the same is true in the case of our Borel local lemma.

(iii) If the set $V(G)$ of vertices is countable, it can be regarded as a standard Borel space when we equip it with the discrete Borel structure. As will turn out, in this situation all local rules on G are Borel. If $V(G)$ is finite then $\mathrm{Rel}(G)$ is also finite, and hence of uniformly subexponential growth. Thus Theorem 1.3 includes Theorem 1.1 as a special case when $V(G)$ is a finite set.

The Moser-Tardos algorithm with limited randomness. The technique we use to prove Theorem 1.3 is a modified Moser-Tardos algorithm, and it is of independent interest.

The Moser-Tardos algorithm (MTA) is a randomized algorithm for *finding* the satisfying assignment under the assumptions of the local lemma. A version of it has been first described by Moser [5], and a modified version has been described by Moser and Tardos [6]. We refer to the introduction of [6] for the history of attempts to find a constructive version of the local lemma.

To motivate our modified MTA let us recall the parallel version of the MTA. Let us assume that the set $V(G)$ of vertices is finite. We start by "sampling" each point of $V(G)$ at random, i.e. we choose uniformly at random a function $f_0 \in b^{V(G)}$. Now we choose a subset $W_0 \subset V(G)$ which is maximal among the subsets of $V(G)$ satisfying the following two properties:

(i) The function f_0 violates the local rule at all points of W_0. More precisely, for all $x \in W_0$ we have that f_0 restricted to the neighbourhood of x is *not* an element of $\mathbf{R}(x)$.
(ii) W_0 is an independent set in the graph $\mathrm{Rel}(G)$.

We define f_1 by "resampling f_0 at variables in W_0". More precisely, we start by defining X_0 to be the union of neighbourhoods of points in W_0, and we let $Y_0 := V(G) \setminus X_0$. Now, we define f_1 to be equal to f_0 on Y_0. Finally, we choose uniformly at random a function in b^{X_0}, and we define f_1 to be equal to that function on X_0.

We repeat this procedure with f_1 in place of f_0, and so on, until we end up with a satisfying assignment. With overwhelming probability a satisfying assignment will be found in a time which is linear in $\log(|V(G)|)$.

Let us informally describe how we modify the MTA. We partition $V(G)$ into p disjoint parts $V_0, \dots, V_{p-1}$ with the property that for every $i \in p$ and all distinct $x, y \in V_i$ we have that x and y are at least r far away from each other in the graph $\mathrm{Rel}(G)$, where the choice of r depends only on the growth of the balls in $\mathrm{Rel}(G)$, but not on $|V(G)|$. Now we assign to each part V_i a "source of randomness", i.e. an element $\mathtt{rnd}$ of $b^{\mathbb{N}}$, and we use this single sequence for resamplings of *all* points x which are in V_i.

Thus the points which lie in the same parts are no longer resampled independently from each other. We refer to this version of the MTA as the *MTA with limited randomness* and we describe it precisely in Section 2.

Previous results and open questions. To our best knowledge, there is no previous work which establishes any *Borel* variants of the local lemma.

However, measurable variants of the local lemma have been studied by Kun [4] and very recently by Bernshteyn [2] (the present work has been carried out independently of [2]). Let us discuss those of the results of [4] and [2] which are related to the present work.

We warn the reader that the following definition is not equivalent to the similar notion in [2].

Definition 1.6. Let Γ be a countable group, let X be a standard Borel space, let ν be a probability measure on X and let $\rho\colon \Gamma \curvearrowright X$ be an action by measure-preserving Borel bijections. For a sequence $\gamma_0, \dots, \gamma_{k-1}$ of elements of Γ we define $G = G(\gamma_0, \dots, \gamma_{k-1})$ to be the graph with $V(G) := X$ and $(x, y) \in E(G)$ if for some $i \in k$ we have $\gamma_i.x = y$.

We say that *the measurable local lemma holds for the action* ρ if for all sequences $\gamma_0, \dots, \gamma_{k-1}$ and all Borel local rules on $G(\gamma_0, \dots, \gamma_{k-1})$ such that $p(x) < \frac{1}{e\Delta}$ for all $x \in X$, there exists a measurable function $f\colon X \to b$ which satisfies $\mathbf{R}$.

The methods which we use to prove Theorem 1.3 can be rather easily modified to prove the following theorem.

Theorem 1.7 (Measurable Local Lemma)**.** *Let Γ be a countable amenable group, let X be a standard Borel space, let ν be a probability measure on X and let $\rho\colon \Gamma \curvearrowright X$ be an action by measure-preserving Borel bijections which is essentially free, i.e. for ν-almost all $x \in X$ we have that the map $\Gamma \to X$ given by $\gamma \mapsto \gamma.x$ is a bijection. Then the measurable local lemma holds for ρ.*

We do not provide a proof of this theorem in the present paper, but it will be included in a future work. However, if Γ is of subexponential growth (hence in particular amenable), then Theorem 1.7 clearly follows from Theorem 1.3.

In [4] it is shown that the standard Moser-Tardos algorithm can be applied in the setting of an infinite countable graph. As a corollary, the measurable local lemma holds for the Bernoulli action $\Gamma \curvearrowright [0,1]^\Gamma$, where Γ is an arbitrary (not necessarily amenable) group . In [2] it is shown that the meaurable local lemma holds for any action $\Gamma \curvearrowright X$ for which there exists a Γ-equivariant Borel surjection onto $[0,1]^\Gamma$.

In particular, if Γ is an amenable group then the results of [4] and [2] imply the measurable local lemma only for the actions $\Gamma \curvearrowright X$ which have infinite entropy. This is a rather large constrain - many natural actions of amenable groups which are covered by Theorems 1.3 and Theorems 1.7 do not have infinte entropy. For example, as far as we know, it is impossible to deduce the measurable, let alone Borel, version of Corollary 1.2 from [4] or [2]

On the other hand, the results in [4] and [2] do not require the group Γ to be amenable. It is very interesting open problem to determine whether the measurable local lemma holds for all probability measure preserving actions of all groups.

Let us make also the following definition.

Definition 1.8. Let Γ be a countable group, let X be a standard Borel space, and let $\rho\colon \Gamma \curvearrowright X$ be an action by Borel bijections. As before , for a sequence $\gamma_0, \dots, \gamma_{k-1}$ of elements of Γ we define $G = G(\gamma_0, \dots, \gamma_{k-1})$ to be the graph with $V(G) := X$ and $(x, y) \in E(G)$ if for some $i \in k$ we have $\gamma_i.x = y$.

We say that *the Borel local lemma holds for the action* ρ if for all sequences $\gamma_0, \dots, \gamma_{k-1}$ and all Borel local rules on $G(\gamma_0, \dots, \gamma_{k-1})$ such that $p(x) < \frac{1}{e\Delta}$ for all $x \in X$, there exists a Borel function $f\colon X \to b$ which satisfies $\mathbf{R}$.

Theorem 1.3 establishes in particular that the Borel local lemma holds for all Borel actions of all groups of subexponential growth. In view of Theorem 1.7 it would be interesting to know whether the Borel local lemma holds for all Borel actions of all amenable groups. In fact, as far as we know, it could be that the Borel local lemma holds for all Borel actions of all groups.

1.a. **Basic notation and definitions.**

Given a set X, we denote by $\mathcal{P}(X)$ the power set of X. Given another set Y we let X^Y be the set of all functions from Y to X. We do not automatically identify the power set of X with 2^X, which is the set of all functions from X to $2 = \{0, 1\}$. The cardinality of X is denoted by $|X|$.

As already mentioned, we frequently use the fact that any natural number is equal to the set of all smaller natural numbers. This occasionally leads to the following notational clash: when $p \in \mathbb{N}$, the symbol 2^p denotes both the set of functions from p to 2, and the cardinality of that set. However, it should not lead to a confusion.

An oriented graph is a pair (V, E) where V is a set and E is a subset of $V \times V$. In particular, oriented graphs are allowed to have self-loops, but each vertex can have at most one self-loop. An oriented graph G is *symmetric* if for any two vertices $x, y \in V(G)$ we have that $(x, y) \in E(G)$ if and only if $(y, x) \in E(G)$. The *symmetrization* of an oriented graph G is the smallest symmetric graph H such that $V(G) = V(H)$ and $E(G) \subset E(H)$.

For $x \in V(G)$ we let $\mathrm{Var}(x) := \{y \in V(G)\colon (x, y) \in E(G)\}$, $\mathrm{Cla}(x) := \{y \in V(G)\colon (y, x) \in E(G)\}$, and $N(x) := \mathrm{Var}(x) \cup \mathrm{Cla}(x)$. If we need to stress the dependence on G we use the notation $\mathrm{Var}_G(x)$, $\mathrm{Cla}_G(x)$, and $N_G(x)$. If A is a set of vertices then we let $\mathrm{Var}(A) := \bigcup_{x \in A} \mathrm{Var}(x)$, and similarly for $\mathrm{Cla}(A)$ and $N(A)$.

We define $\mathrm{Rel}(G)$ to be the oriented graph on the same set of vertices as G and with an oriented edge from x to y if and only if $\mathrm{Var}_G(x) \cap \mathrm{Var}_G(y)$ is non-empty. Clearly $\mathrm{Rel}(G)$ is symmetric.

A *variable graph* is an oriented graph G together with two families $a(x)$, $x \in V(X)$, and $b(x)$, $x \in V(X)$, where $a(x)$ is a well-order on $\mathrm{Var}(x)$ and $b(x)$ is a well-order on $\mathrm{Cla}(x)$.

A *relation graph* is an oriented graph with edges labelled by natural numbers, in such a way that at each vertex x all out-edges starting at x have different labels.

If G is a variable graph then we consider $\mathrm{Rel}(G)$ to have the following relation graph structure. For every vertex x we label the edges starting at x with distinct elements of $|N_{\mathrm{Rel}(G)}(x)|$, which we do by ordering the elements of $N_{\mathrm{Rel}(G)}(x)$ in the following way. Let $y, z \in N_{\mathrm{Rel}(G)}(x)$ be two distinct elements, let v and w be the smallest elements of, respectively, $\mathrm{Var}_G(x) \cap \mathrm{Var}_G(y)$ and $\mathrm{Var}_G(x) \cap \mathrm{Var}_G(z)$, with respect to the order on

$\mathrm{Var}_G(x)$. If v is smaller than w in $\mathrm{Var}_G(x)$ then we define $y < z$. If w is smaller than v then we define $z < y$. If $v = w$, then we define $y < z$ if and only if y is smaller than z in $\mathrm{Cla}_G(v)$.

Example 1.9. Let us explain how to obtain a variable graph from an instance of 3-SAT in conjunctive normal form (CNF). An instance of 3-SAT in CNF consists of a sequence $x_0, x_1, \dots, x_{n-1}$ of pair-wise distinct variables and a sequence $c_0, c_1, \dots, c_{m-1}$ of pair-wise distinct *clauses*, i.e. each c_i is of the form $a_0 \vee a_1 \vee a_2$, where each a_j is equal either to some x_k or $\neg x_k$. We say that such an instance is *satisfiable* if there exists a function $f\colon \{x_0, \dots, x_{n-1}\} \to \{0, 1\}$ such that all the clauses c_i are satisfied.

We associate to such an instance a variable graph G as follows. We let $V(G) := \{c_0, \dots, c_{m-1}\} \sqcup \{x_0, \dots, x_{n-1}\}$, and we let (c_i, x_j) be an edge if either x_j or $\neg x_j$ appears in c_i. In particular the graph G is bipartite and every edge is oriented from some c_i towards some x_j. The well-orders necessary to finish the definition of a variable graph are induced from the orders $x_0 < x_1 < \ldots < x_{n-1}$ and $c_0 < c_1 < \ldots < c_{m-1}$. It is easy to check that $\mathrm{Rel}(G)$ is a graph on the same set of vertices, such that all edges are of the form (c_i, c_j) for some clauses c_i and c_j, and (c_i, c_j) is an edge if c_i and c_j share a variable. In particular every vertex c_i has a self-loop.

2. Moser-Tardos algorithm

2.a. **Borel graphs.**

Let Ω be a standard Borel space and $\mathcal{B}$ be its Borel σ-algebra. A *Borel arrow on* Ω is a pair (U, γ), where $U \in \mathcal{B}$ and $\gamma\colon U \to \Omega$ is a Borel isomorphism onto its image. Given a countable sequence $S = ((U_1, \gamma_1), (U_2, \gamma_2), \dots)$ of Borel arrows, we associate to it an oriented graph $\mathcal{G}(S)$ as follows. We let $V(\mathcal{G}(S)) := \Omega$ and

$$E(\mathcal{G}(S)) := \{(x, y)\colon \exists i \in \mathbb{N} \text{ such that } x \in U_i \text{ and } \gamma_i(x) = y\}.$$

A *Borel graph* is a graph $\mathcal{G}$ which is equal to $\mathcal{G}(S)$ for some sequence S of Borel arrows on a standard Borel space.

The following properties are easy to check.

Lemma 2.1. *Let $\mathcal{G}$ be a Borel graph. We have the following properties.*

(i) The set $E(\mathcal{G})$ of edges of $\mathcal{G}$ is a Borel subset of $V(\mathcal{G}) \times V(\mathcal{G})$.
(ii) If $X \subset \Omega$ is a Borel set then the graph induced on X from $\mathcal{G}(S)$ is also a Borel graph.
(iii) The symmmetrization of $\mathcal{G}$ is a Borel graph.
(iv) The graph $\mathrm{Rel}(\mathcal{G})$ is a Borel graph.
(v) If $A \subset V(\mathcal{G})$ is a Borel set then also the sets $\mathrm{Var}(A)$, $Cl(A)$ and $N(A)$ are Borel. □

An *independence function for a Borel graph* $\mathcal{G}$ is a function $\mathrm{I}\colon \mathcal{B} \to \mathcal{B}$ with the property that for every $X \in \mathcal{B}$ the set $\mathrm{I}(X)$ is a maximal subset of X which is independent in $\mathcal{G}$.

In the appendix we show that every locally finite Borel graph admits an independence function.

Let $S = ((U_1, \gamma_1), (U_2, \gamma_2), \dots)$ be a sequence of Borel arrows on a standard Borel space. The Borel graph $\mathcal{G}(S)$ becomes a variable graph when we induce the orders on $\mathrm{Var}(x)$ and $\mathrm{Cla}(x)$ from the sequence S in the following way. For two distinct vertices $y, z \in \mathrm{Var}(x)$ we let $y < z$ if and only if there exists i such that $y = \gamma_i(x)$ and for all $j < i$ we have $z \neq \gamma_j(x)$. Similarly for two distinct vertices $y, z \in \mathrm{Cla}(x)$ we let $y < z$ if and only if there exists i such that $x = \gamma_i(y)$ and for all $j < i$ we have $x \neq \gamma_j(z)$. A *Borel variable graph* is a Borel graph with the structure of a variable graph just described.

2.b. **Moser-Tardos algorithm with limited randomness.**

It is convenient to fix a natural number $b > 1$ for the rest of this article.

A *local rule* for a Borel variable graph $\mathcal{G}$ is a map $\mathbf{R}$ whose domain is $V(\mathcal{G})$ and which assigns to each $x \in V(G)$ a subset $\mathbf{R}(x)$ of $b^{\mathrm{Var}(x)}$.

Let $\mathbb{S}$ denote the set of all finite b-valued sequences. In other words, the set $\mathbb{S}$ consists of all functions $f\colon k \to b$, where $k \in \mathbb{N}$. Since for every $x \in V(\mathcal{G})$ the set $\mathrm{Var}(x)$ is well-ordered, we have a canonical bijection $\iota_x\colon |\mathrm{Var}(x)| \to \mathrm{Var}(x)$. Thus a local rule can be canonically identified with a function $\mathbf{R}'\colon V(\mathcal{G}) \to \mathcal{P}(\mathbb{S})$, as follows: we let $\mathbf{R}'(x)$ be the set $\{f \circ \iota_x\colon f \in \mathbf{R}(x)\}$.

We say that $\mathbf{R}$ is *Borel* if the preimage of every element of $\mathcal{P}(\mathbb{S})$ under $\mathbf{R}'$ is Borel.[1] The only purpose of defining the function $\mathbf{R}'$ was to define what it means for $\mathbf{R}$ to be Borel, and in fact $\mathbf{R}'$ will not be used again.

Given a function $f\colon V(\mathcal{G}) \to b$ and $x \in V(\mathcal{G})$ we let $\overline{f}(x)\colon \mathrm{Var}(x) \to b$ be the restriction of f to $\mathrm{Var}(x)$. We say that f *satisfies* $\mathbf{R}$ if for every x we have $\overline{f}(x) \in \mathbf{R}(x)$.

Given $\mathbf{R}$, the *complementary rule* $\mathbf{R}^c$ is defined by setting $\mathbf{R}^c(x) := b^{\mathrm{Var}(x)} \setminus \mathbf{R}(x)$. For $f\colon V(\mathcal{G}) \to b$ we set $\mathrm{B}_{\mathbf{R}}(f) := \{x \in V(\mathcal{G})\colon \overline{f}(x) \in \mathbf{R}^c(x)\}$. When $\mathbf{R}$ is clear from the context we write B instead of $\mathrm{B}_{\mathbf{R}}$.

The proof of the following lemma is a tedious but routine calculus exercise, and so we state it without a proof.

Lemma 2.2. *Let $\mathcal{G}$ be a locally finite Borel variable graph. If $f\colon V(\mathcal{G}) \to b$ is a Borel function and $\mathbf{R}$ is a Borel local rule for $\mathcal{G}$ then $\mathrm{B}_{\mathbf{R}}(f)$ is a Borel subset of $V(\mathcal{G})$.* □

To avoid cluttering we will write $\mathrm{I}\,\mathrm{B}(f)$ instead of $\mathrm{I}(\mathrm{B}(f))$ throughout the article.

A *Borel partition* of $V(\mathcal{G})$ is a tuple $\pi = (W_0, \dots, W_{p-1})$ such that for all $i \in p$ we have that W_i is a Borel subset of $V(\mathcal{G})$, the sets W_i are pair-wise disjoint, and $\bigcup_{i \in p} W_i = V(\mathcal{G})$. Given a Borel partition $\pi = (W_0, \dots, W_{p-1})$ and $x \in V(\mathcal{G})$ we define $\pi(x)$ to be the unique $i \in p$ such that $x \in W_i$.

[1] This definition makes sense for all Borel variable graphs $\mathcal{G}$. However, it is a reasonable definition only when $\mathcal{G}$ is locally finite, in which case the image of $\mathbf{R}'$ is countable.

A *Moser-Tardos tuple* is a tuple $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi)$, where $\mathcal{G}$ is a Borel variable graph, $\mathbf{R}$ is a Borel local rule for $\mathcal{G}$, I is an independence function for $\mathrm{Rel}(\mathcal{G})$, and $\pi = (W_0, W_1, \dots, W_{p-1})$ is a Borel partition of $V(\mathcal{G})$.

We are ready to present a version of the Moser-Tardos algorithm. Let $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi)$ be a Moser-Tardos tuple, let $k \in \mathbb{N}$ and $\mathtt{rnd} \in b^{p \times k}$. The *$k$-step Moser-Tardos algorithm for* $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi)$ and $\mathtt{rnd}$ takes as its input a Borel function $f \colon V(\mathcal{G}) \to b$ and returns a sequence of Borel functions $\mathrm{M}_0 := f, \mathrm{M}_1, \dots, \mathrm{M}_k \colon V(\mathcal{G}) \to b$. Furthermore if $k' > k$ and $\mathtt{rnd}' \in b^{p \times k'}$ extends $\mathtt{rnd}$ in the obvious sense, then the k'-step Moser-Tardos algorithm extends the k-step Moser-Tardos algorithm in the obvious sense. Thus we will also speak of the *$\mathbb{N}$-step Moser-Tardos algorithm* which outputs an infinite sequence of functions $\mathrm{M}_0, \mathrm{M}_1, \mathrm{M}_2, \dots$.

It is convenient to fix $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi)$ and f, and regard the Moser-Tardos algorithm as only depending on $\mathtt{rnd}$. As such, we do not incorporate $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi)$ and f in the notation, and we include $\mathtt{rnd}$ in the notation only when it is necessary to stress the dependence on $\mathtt{rnd}$. The function $\mathtt{rnd}$ should be thought of as the source of randomness for the Moser-Tardos algorithm (see also Remark 2.3 below). The tuple $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, f)$ will be called a *long Moser-Tardos tuple*.

We let $\mathrm{M}_0 := f$ and we proceed to define $\mathrm{M}_1 = \mathrm{M}_1(\mathtt{rnd}), \dots, \mathrm{M}_k = \mathrm{M}_k(\mathtt{rnd})$ inductively.

Let $h_0 \colon V(\mathcal{G}) \to \mathbb{N}$ be identically equal to 0. Now suppose that for some $j < k$ the functions M_j and h_j are defined. We let

$$h_{j+1}(x) = \begin{cases} h_j(x) + 1 & \text{if } x \in \mathrm{Var}(\mathrm{I}\,\mathrm{B}(\mathrm{M}_j)), \\ h_j(x) & \text{otherwise,} \end{cases}$$

and

$$\mathrm{M}_{j+1}(x) = \begin{cases} \mathtt{rnd}(\pi(x), h_j(x)) & \text{if } x \in \mathrm{Var}(\mathrm{I}\,\mathrm{B}(\mathrm{M}_j)), \\ \mathrm{M}_j(x) & \text{otherwise.} \end{cases}$$

Using Lemmas 2.2 and 2.1(v), it is easy to prove inductively that the functions $\mathrm{M}_0, \dots, \mathrm{M}_k$ are Borel.

If for some $x \in V(\mathcal{G})$ we have $h_k(x) > N$ then we say that the k-step Moser-Tardos algorithm *resamples x more than N times*. Similarly if for some $k \in \mathbb{N}_+$ and $x \in V(\mathcal{G})$ we have $h_k(x) > N$ then we say that the $\mathbb{N}$-step Moser-Tardos algorithm *resamples x more than N times*.

If $N \in \mathbb{N}$ is such that for all $x \in V(\mathcal{G})$ we have $h_k(x) \leqslant N$ then we say that the k-step Moser-Tardos algorithm *resamples each point at most N times*. Similarly when $N \in \mathbb{N}$ is such that for all $k \in \mathbb{N}$ and all $x \in V(\mathcal{G})$ we have $h_k(x) \leqslant N$ then we say that the $\mathbb{N}$-step Moser-Tardos algorithm *resamples each point at most N times*.

Remark 2.3. Throughout the whole article, whenever we speak of the probability of an event, the implied probability space is one of the spaces $b^{p \times k}$, $k \in \mathbb{N}$, or $b^{p \times \mathbb{N}}$, together with the product measure of the uniform measures on the copies of b. Thus for example the statement "With probability η the $\mathbb{N}$-step Moser-Tardos algorithm resamples each point

at most N times" should be read as: "Let $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi)$ and f be fixed. The measure of the set of those $\mathtt{rnd} \in b^{p\times\mathbb{N}}$ such that the $\mathbb{N}$-step Moser-Tardos algorithm for $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, \mathtt{rnd})$ with input f resamples each point at most N times is equal to η".

Finally we define two functions, Used and Unused, which will be useful in the upcoming analysis of the k-step Moser-Tardos algorithm. Both Used and Unused assign a b-valued sequence to each vertex $x \in V(\mathcal{G})$. We define $\textsc{Used}(x)$ to be the sequence

$$\mathtt{rnd}(\pi(x), 0), \dots, \mathtt{rnd}(\pi(x), h_{k-1}(x) - 1),$$

and $\textsc{Unused}(x)$ to be the sequence

$$\mathtt{rnd}(\pi(x), h_{k-1}(x)), \dots, \mathtt{rnd}(\pi(x), k - 1).$$

When we need to stress the dependence of $h_k(x)$, $\textsc{Used}(x)$ and $\textsc{Unused}(x)$ on $\mathtt{rnd}$, we write $h_k(\mathtt{rnd}, x)$, $\textsc{Used}(\mathtt{rnd}, x)$ and $\textsc{Unused}(\mathtt{rnd}, x)$, respectively.

Let us set $X(\mathcal{G}, \mathbf{R}) := \{x \in V(\mathcal{G})\colon \mathbf{R}(x) \neq b^{\mathrm{Var}_{\mathcal{G}}(x)}\}$. We will need to deal with the set $X(\mathcal{G}, \mathbf{R})$ only to reduce the general case of the local lemma to the case when for all $x \in V(\mathcal{G})$ the sets $|Var(x)|$ have the same number of elements.

We finish with the following simple observation.

Lemma 2.4. *Let $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, f)$ and $(\overline{\mathcal{G}}, \overline{\mathbf{R}}, \overline{\mathrm{I}}, \overline{\pi}, \overline{f})$ be two long Moser-Tardos tuples, such that*

(i) $X(\overline{\mathcal{G}}, \overline{\mathbf{R}}) \subset V(\mathcal{G})$, and $\mathcal{G}$ is an induced subgraph of $\overline{\mathcal{G}}$,
(ii) for all $x \in V(\mathcal{G})$ we have $\pi(x) = \overline{\pi}(x)$ and $f(x) = \overline{f}(x)$,
(iii) for all $x \in V(\mathcal{G})$ and all $g \in b^{\mathrm{Var}_{\mathcal{G}}(x)}$ we have that $g \in \mathbf{R}(x)$ if and only if g is a restriction of an element of $\overline{\mathbf{R}}(x)$,
(iv) for all Borel subsets $U \subset V(\mathcal{G})$ we have $\mathrm{I}(U) = \overline{\mathrm{I}}(U)$, and
(v) for $x \in V(\overline{\mathcal{G}}) \setminus V(\mathcal{G})$ we have $|\operatorname{Cla}_{\overline{\mathcal{G}}}(x) \cap V(\mathcal{G})| \leqslant 1$.

Let $k, N \in \mathbb{N}$. The probability that the k-step Moser-Tardos algorithm for $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, f)$ resamples each point at most N times is equal to the probability that the k-step Moser-Tardos algorithm for $(\overline{\mathcal{G}}, \overline{\mathbf{R}}, \overline{\mathrm{I}}, \overline{\pi}, \overline{f})$ resamples each point at most N times.

Proof. Let p and $\bar{p}$ denote respectively the numbers of parts of the partitions π and $\overline{\pi}$. Let $\mathtt{rnd} \in 2^{p\times k}$ and let $\overline{\mathtt{rnd}} \in 2^{\bar{p}\times k}$ be an extension of $\mathtt{rnd}$, i.e. for $i \in p$ and $j \in k$ we have $\mathtt{rnd}(i, j) = \overline{\mathtt{rnd}}(i, j)$.

Let $\mathrm{M}_0, \dots, \mathrm{M}_k$ and $\overline{\mathrm{M}}_0, \dots, \overline{\mathrm{M}}_k$ be the sequence of functions produced by the respective Moser-Tardos algorithms. It is clear that for all $i \in k$ and $x \in V(\mathcal{G})$ we have $\mathrm{M}_i(x) = \overline{\mathrm{M}}_i(x)$. It is also clear that $\mathrm{B}_{\overline{\mathbf{R}}}(\overline{\mathrm{M}}_i)$ is contained in $V(\mathcal{G})$. Thus, we have $\overline{\mathrm{I}}\,\mathrm{B}_{\overline{\mathbf{R}}}(\overline{\mathrm{M}}_i) = \mathrm{I}\,\mathrm{B}_{\mathbf{R}}(\mathrm{M}_i)$.

In particular any point $x \in V(\mathcal{G})$ is resampled exactly the same amount of times by the Moser Tardos agorithms for $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, f)$ and for $(\overline{\mathcal{G}}, \overline{\mathbf{R}}, \overline{\mathrm{I}}, \overline{\pi}, \overline{f})$.

For $x \in V(\overline{\mathcal{G}}) \setminus V(\mathcal{G})$ we consider two cases. If $\mathrm{Cla}_{\overline{\mathcal{G}}}(x) \cap V(\mathcal{G})$ is empty then the Moser-Tardos algorithm on $(\overline{\mathcal{G}}, \overline{\mathbf{R}}, \overline{\mathrm{I}}, \overline{\pi}, \overline{f})$ does not resample x at all. If $\mathrm{Cla}_{\overline{\mathcal{G}}}(x) \cap V(\mathcal{G})$ is not empty then it contains a unique point $y \in V(\mathcal{G})$. If there is $z \in \mathrm{Var}_{\mathcal{G}}(y)$, then clearly x is resampled exactly as many times as z. Finally if $\mathrm{Var}_{\mathcal{G}}(y)$ is empty then clearly for any $F\colon V(\overline{\mathcal{G}}) \to b$ we have $y \notin \mathrm{B}(F)$, and hence again the Moser-Tardos algorithm does not resample x at all. This finishes the proof. □

3. Preliminaries on landscapes of trees

3.a. **Landscapes.**

For $x \in V(G) \times \mathbb{N}$ we let $\overline{x}$ be the first coordinate of x and we let the *level of* x, denoted by $\ell(x)$, be the second coordinate of x. For $i \in \mathbb{N}$ and $V \subset V(G) \times \mathbb{N}$ we let V_i be the subset of those $x \in V$ for which $\ell(x) = i$.

Given a variable graph G we define a relation graph $\mathrm{Canvas}(G)$ as follows. The set of vertices of $\mathrm{Canvas}(G)$ is $V(G) \times \mathbb{N}$. For every edge (x, y) of the graph $\mathrm{Rel}(G)$ and every $i \in \mathbb{N}$ we add an edge $((x, i), (y, i+1))$ to $\mathrm{Canvas}(G)$, with the same label as the label of (x, y).

A *pseudo-landscape* $\mathcal{L}$ is a triple consisting of the following elements.

(i) a variable graph $G_{\mathcal{L}}$,
(ii) a local rule $\mathbf{R}_{\mathcal{L}}$ on $G_{\mathcal{L}}$,
(iii) a subgraph $\mathrm{For}(\mathcal{L})$ of $\mathrm{Canvas}(G_{\mathcal{L}})$ such that at each vertex the in-degree is either 0 or 1.

Note that the graph in (iii) is in fact a forest and it will be referred to as *the forest of* $\mathcal{L}$. We let $V(\mathcal{L})$ and $E(\mathcal{L})$ denote respectively the sets of vertices and edges of $\mathrm{For}(\mathcal{L})$, and we let $\mathrm{T}(\mathcal{L})$ be the set of connected components of $\mathrm{For}(\mathcal{L})$. For $\tau \in \mathrm{T}(\mathcal{L})$ we let $\rho(\tau) \in V(\tau)$ be the root, i.e. the unique vertex with minimal level. Finally, we let $\ell(\tau)$ be the *level of* τ, which is by definition equal to $\ell(\rho(\tau))$.

We let $\mathrm{Canvas}(G_{\mathcal{L}}, k)$ to be the subgraph of $\mathrm{Canvas}(G_{\mathcal{L}})$ induced on the set $V(G_{\mathcal{L}}) \times k$ of vertices. The minimal k such that the forest in (iii) is a subgraph of $\mathrm{Canvas}(G_{\mathcal{L}}, k)$ is called the *height of* $\mathcal{L}$.

A pseudo-landscape is a *landscape* if for all $i \in \mathbb{N}$ and all distinct $x, y \in V(\mathcal{L})_i$ we have that the distance between $\overline{x}$ and $\overline{y}$ in $\mathrm{Rel}(G_{\mathcal{L}})$ is at least 2.

Let $D, \Delta, \beta \in \mathbb{N}$. We say that $\mathcal{L}$ is of of type (D, Δ, β) if the maximal out-degree in $G_{\mathcal{L}}$ is at most D, the maximal degree in $\mathrm{Rel}(G_{\mathcal{L}})$ is at most Δ, and for every $x \in V(G_{\mathcal{L}})$ we have $|\mathbf{R}^c_{\mathcal{L}}(x)| \leqslant \beta$.

We finish by defining a *decoration* of a landscape $\mathcal{L}$ as consisting of the following data.

(1) a function $\textsc{Final}(\mathcal{L}) \in b^{V(G_{\mathcal{L}})}$
(2) for each $x \in V(\mathcal{L})$ an element $\textsc{Prev}(x)$ of $\mathbf{R}^c_{\mathcal{L}}(\overline{x})$. In particular $\textsc{Prev}(x)$ is a function $\mathrm{Var}(\overline{x}) \to b$.

(3) a function $\pi\colon V(G_{\mathcal{L}}) \to \mathbb{N}$.

Let $D, \Delta, \beta,\ N_1, N_2, p \in \mathbb{N}$. We say that a decorated landscape $\mathcal{L}$ is of type $(D, \Delta, \beta,\ N_1, N_2, p)$ if it is of type (D, Δ, β) and additionally $|V(G_{\mathcal{L}})| \leqslant N_1$, $|V(\mathcal{L})| = N_2$, and $\mathrm{im}(\pi) \subset p$.

3.b. **Grounded landscapes.**

Let $\mathcal{L}$ be a landscape. We say that $\mathcal{L}$ is *grounded* if all trees of $\mathcal{L}$ are at level 0. The purpose of this subsection is to define an equivalence relation on landscapes and to show that every landscape is equivalent to a grounded one. We start with several definitions.

We say that $\mathcal{L}$ is *tight* if at least one tree is at level 0 and for every tree $\tau \in \mathrm{T}(\mathcal{L})$ such that $\ell(\tau) > 0$ there exist $\sigma \in \mathrm{T}(\mathcal{L})$ different than τ, $x \in V(\sigma)$ and $y \in V(\tau)$ such that (x, y) is an edge of $\mathrm{Canvas}(G_{\mathcal{L}})$. Note that in particular any grounded landscape is tight.

For a vertex x in $\mathrm{Canvas}(G_{\mathcal{L}})$ such that $\ell(x) > 0$ we let $\mathrm{Push}(x) := (\overline{x}, \ell(x) - 1)$. If G is a subgraph of $\mathrm{Canvas}(G_{\mathcal{L}})$ such that for all $x \in V(G)$ we have $\ell(x) > 0$ then we let $\mathrm{Push}(G)$ be the graph whose set of vertices is $\{\mathrm{Push}(x)\colon x \in V(G)\}$ and whose set of edges is $\{(\mathrm{Push}(x), \mathrm{Push}(y))\colon (x, y) \in E(G)\}$.

If for all $x \in V(\mathcal{L})$ we have $\ell(x) > 0$ then we say that $\mathcal{L}$ is *pushable* and we define $\mathrm{Push}(\mathcal{L})$ to be the triple consisting of the graph $G_{\mathcal{L}}$, the local rule $\mathbf{R}_{\mathcal{L}}$, and the graph $\mathrm{Push}(\mathrm{For}(\mathcal{L}))$.

For $\tau \in \mathrm{T}(\mathcal{L})$ such that $\ell(\tau) > 0$ we let $\mathrm{Push}(\mathcal{L}, \tau)$ be the triple consisting of the graph $G_{\mathcal{L}}$, the local rule $\mathbf{R}_{\mathcal{L}}$ and the subgraph of $\mathrm{Canvas}(G_{\mathcal{L}})$ which is the union of the graph $\mathrm{Push}(\tau)$ and of all the graphs in $\mathrm{T}(\mathcal{L}) \setminus \{\tau\}$.

We say that $\tau \in \mathrm{T}(\mathcal{L})$ is *pushable* if τ is a tree which "witnesses that $\mathcal{L}$ is not tight", i.e. such that $\ell(\tau) > 0$ and for all $\sigma \in \mathcal{L}$ different than τ, and all $x \in V(\sigma)$, $y \in V(\tau)$ we have that (x, y) is not an edge in $\mathrm{Canvas}(G_{\mathcal{L}})$.

The landscapes $\mathrm{Push}(\mathcal{L})$ for a pushable $\mathcal{L}$ and $\mathrm{Push}(\mathcal{L}, \tau)$, for a pushable τ will be referred to as *simple pushes*. The following is straightforward to check.

Lemma 3.1. *A simple push of a landscape is a landscape.* □

Let (x, y, z) be a triple of distinct elements of $V(\mathcal{L})$. We say (x, y, z) is *rebranchable* if $(x, z) \in E(\mathcal{L})$ and (y, z) is an edge of $\mathrm{Canvas}(G_{\mathcal{L}})$.

A *rebranching* of $\mathcal{L}$ with respect to a rebranchable triple (x, y, z) is a triple consisting of $G_{\mathcal{L}}$, $\mathbf{R}_{\mathcal{L}}$ and the subgraph H of $\mathrm{Canvas}(G_{\mathcal{L}})$ defined as follows. We let $V(H) := V(\mathcal{L})$ and $E(H) = E(\mathcal{L}) \setminus \{(x, z)\} \cup \{(y, z)\}$. A tree $\tau \in \mathrm{T}(\mathcal{L})$ is *rebranchable* if there exists a rebranchable triple (x, y, z) such that $(x, z) \in E(\tau)$.

We say that a pair (y, z) of distinct elements of $V(\mathcal{L})$ is *joinable* if z is a root of a tree in $\mathrm{T}(\mathcal{L})$ and (y, z) is an edge in $\mathrm{Canvas}(G_{\mathcal{L}})$. A tree $\tau \in \mathrm{T}(\mathcal{L})$ is *joinable* if for some $y \in V(\mathcal{L})$ the pair (y, $\rho(\tau)$) is joinable.

A *joining* of $\mathcal{L}$ with respect to a joinable pair (y, z) is a triple which arises from $\mathcal{L}$ by adding the edge (y, z) to $\mathrm{For}(\mathcal{L})$.

The following lemma is straightforward to check.

Lemma 3.2. *Let $\mathcal{L}$ be a landscape.*

(i) If (x, y, z) is rebranchable triple then the rebranching of $\mathcal{L}$ with respect to (x, y, z) is a landscape.
(ii) If (y, z) is a joinable pair then the joining of $\mathcal{L}$ with respect to (y, z) is a landscape.
(iii) $\mathcal{L}$ is not tight if and only if $\mathcal{L}$ is pushable or there exists $\tau \in \mathrm{T}(\mathcal{L})$ which is pushable.
(iv) A tree $\tau \in \mathrm{T}(\mathcal{L})$ is pushable if and only if $\ell(\tau) > 0$, it is not rebranchable and it is not joinable. □

We say that two landscapes are *equivalent* if one arises from the other by a sequence consisting of simple pushes, rebranchings and joinings.

Lemma 3.3. *Every landscape is equivalent to a grounded landscape.*

Proof. For a landscape $\mathcal{L}$ we let $\mathrm{G}(\mathcal{L}) \subset T(\mathcal{L})$ denote the set of trees at level zero. We prove the lemma by showing that for every landscape $\mathcal{L}$ such that $|T(\mathcal{L})| - |\mathrm{G}(\mathcal{L})| > 0$ there exists a sequence of simple pushes, rebranchings, and joinings which produces a landscape $\mathcal{K}$ for which $|T(\mathcal{K})| - |\mathrm{G}(\mathcal{K})| < |T(\mathcal{L})| - |\mathrm{G}(\mathcal{L})|$.

Since a simple push does not change $|T(\mathcal{L})|$ and does not decrease $|G(\mathcal{L})|$, we may assume that $\mathcal{L}$ is tight. Suppose that $\tau \in T(\mathcal{L})$ is such that $\ell(\tau) > 0$.

Since rebranching does not change $|T(\mathcal{L})|$ nor $|\mathrm{G}(\mathcal{L})|$ we may assume that τ is not rebranchable. Thus, by Lemma 3.2, we may assume that τ is joinable. However, performing a joining decreases $|T(\mathcal{L})|$ by 1 and leaves $|G(\mathcal{L})|$ unchanged which finishes the proof. □

We finish by noting that if $\mathcal{L}$ is a decorated landscape, then simple pushes, rebranchings and joinings of $\mathcal{L}$ inherit the decoration from $\mathcal{L}$ in the obvious way.

3.c. **Counting grounded landscapes.**

There is an obvious notion of the *isomorphism of landscapes* $\mathcal{K}$ and $\mathcal{L}$: it is a bijection $V(G_{\mathcal{K}}) \to V(G_{\mathcal{L}})$ which induces an isomorphism of the graphs $G_{\mathcal{K}}$ and $G_{\mathcal{L}}$ compatible with the orders on $\mathrm{Var}(x)$ and $\mathrm{Cla}(x)$ and the local rules, and which induces an isomorphism of $\mathrm{For}(\mathcal{K})$ and $\mathrm{For}(\mathcal{L})$. For decorated landscapes we additionally demand that the induced isomorphisms should be compatible with the decorations in the obvious sense.

In this subsection we will count the iso-classes (i.e. isomorphism classes) of grounded decorated landscapes of a given type. We start by counting subtrees of $\mathrm{Canvas}(G_{\mathcal{L}})$. A Δ-*labelled tree* is an oriented tree such that at each vertex the in-degree is either 0 or 1, and such that the edges are labelled by the elements of Δ in such a way that at each vertex all the out-going edges have different labels.

Lemma 3.4. *Let $\Delta \geqslant 2$. The number of Δ-labeled trees with N vertices is bounded by*

$$\left(\frac{\Delta^\Delta}{(\Delta-1)^{\Delta-1}}\right)^N.$$

Proof. This proof is based on the arguments from [8]. Let $P_0 := 0$ and for a natural number $i > 0$ let P_i be the number of iso-classes of Δ-labelled trees with i vertices. Let

$$P(X) := \sum_{i=0}^{\infty} P_i \cdot X^i,$$

and let ρ be the radius of convergence of $P(X)$. Clearly it is enough to show that $P(\rho) < 1$ and

$$\rho \geqslant \frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}.$$

In fact we will show $P(\rho) \leqslant \frac{1}{\Delta-1}$

Note that in a finite Δ-labelled tree there is a unique vertex with in-degree 0, and we call this vertex the *root*. By considering all the possibilities for the outgoing edges at the root, we arrive at the following equation:

$$P(X) = X(1+P(X))^\Delta.$$

We note that there is only one formal power series $P(X)$ with $P(0) = 0$ which fulfils the above equation.

For $i, j \in \mathbb{N}$ let us define the numbers $Q(i,j) \in \mathbb{N}$ as follows. For all $i \in \mathbb{N}$ we let $Q(i,0) := 0$, and for all $j > 0$ we let $Q(i,j)$ to be the number of Δ-labelled trees with j vertices and such that all vertices are at the distance at most i from the root. Let

$$Q_i(X) = \sum_{j=0}^{\infty} Q(i,j) X^j.$$

The following claim is clear from the definitions.

Claim 3.5. *The polynomials Q_i have non-negative real coefficients, and for every n there exists N such that for all $M > N$ we have that the first n coefficients of $Q_M(X)$ are equal to the first n coefficients of $P(X)$.* □

As a direct corollary we obtain the following claim.

Claim 3.6. *If for some $x \in \mathbb{R}$ the sequence $Q_0(x), Q_1(x), \ldots$ converges to $y \in \mathbb{R}$, then the power series $P(X)$ converges at x and we have $P(x) = y$.*

We proceed to use the above claim to establish the convergence of $P(X)$ for $X \leqslant \frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}$. It is clear that for any i the sequence $Q(i,0), Q(i,1), \ldots$ is non-decreasing. Hence we see that for any $x \geqslant 0$ we have $Q_0(x) \leqslant Q_1(x) \leqslant Q_2(x) \leqslant \ldots$. Also, it is easy to see that each polynomial $Q_i(X)$ is non-decreasing for $X \geqslant 0$. It follows that it is enough to show that the sequence $Q_0(\frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}), Q_1(\frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}), \ldots$ is bounded from above by $\frac{1}{\Delta-1}$.

This is clear for $Q_0(\frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta})$, since $Q_0(X) = X$. Let us prove by induction that for all $i \in \mathbb{N}$ we have $Q_i(\frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}) \leqslant \frac{1}{\Delta-1}$.

Just like in the case of $P(X)$, by considering all the possibilities for the outgoing edges at the root, we see that for all $i \in \mathbb{N}$ we have

$$Q_{i+1}(X) = X(1 + Q_i(X))^\Delta.$$

Thus, using the inductive assumption, we have

$$Q_{i+1}\left(\frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}\right) \leqslant \left(\frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}\right)\left(1 + \frac{1}{\Delta-1}\right)^\Delta = \frac{1}{\Delta-1},$$

which finishes the proof.

□

Lemma 3.7. *Let $D, \Delta, \beta, N_1, N_2, p \in \mathbb{N}$ and $\Delta \geqslant 2$. The number of iso-classes of grounded decorated landscapes $\mathcal{L}$ of type $(D, \Delta, \beta, N_1, N_2, p)$ is at most*

$$C \cdot N_2^{N_1} \cdot \left(\frac{\Delta^\Delta}{(\Delta-1)^{\Delta-1}} \cdot \beta\right)^{N_2},$$

where C depends only on (D, Δ, N_1, p).

Proof. In fact we will show that the number in question is at most

$$(1)\quad N_1 \cdot (N_1+1)^{D \cdot N_1} \cdot (D!)^{N_1} \cdot (\Delta!)^{N_1} \cdot 2^{b^D \cdot N_1} \cdot b^{N_1} \cdot p^{N_1} \cdot N_1^{N_1} \cdot N_2^{N_1} \cdot \left(\frac{\Delta^\Delta}{(\Delta-1)^{\Delta-1}} \cdot \beta\right)^{N_2}.$$

Let us explain all the terms in (1). To define a pseudo-landscape of type $(D, \Delta, \beta, N_1, N_2, p)$ we need to define the graph $G_\mathcal{L}$ of max out-degree at most D on at most N_1 vertices. There are at most $N_1 \cdot (N_1+1)^{D \cdot N_1}$ iso-classes of such graphs, because we can define a graph with vertex set N_1 and for each element of N_1 we give a sequence of D elements, each of which is either an element of N_1 or a special symbol signifying that we put no edge (this special symbol is the reason we have $N_1 + 1$ and not simply N_1). We need the additional factor N_1 to specify which of the vertex sets $1, 2, \ldots, N_1$ is the vertex set of $G_\mathcal{L}$.

The variable graph structure, i.e. the choice of orderings of all the sets $\mathrm{Var}(x)$ and $\mathrm{Cla}(x)$, $x \in V(G_\mathcal{L})$, contributes the factor $(D!)^{N_1} \cdot (\Delta!)^{N_1}$.

We need a local rule $\mathbf{R}_\mathcal{L}$ on $G_\mathcal{L}$. If we fix the graph $G_\mathcal{L}$ with vertex set N_1 then there are at most $2^{b^D \cdot N_1}$ different local rules, because for each vertex $x \in N_1$ we need to specify a family of function in $b^{\mathrm{Var}(x)}$, and there are at most 2^{b^D} such families.

We need to specify the forest of $\mathcal{L}$, and we will do it by specifying a sequence of Δ-labelled trees $T(0), T(1), \ldots, T(N_1 - 1)$, such that $\sum_{i=0}^{N_1} |V(T(i))| = N_2$. We allow some o the trees $T(i)$ to have empty vertex sets. By basic enumerative combinatorics and Lemma 3.4

there are at most

$$N_2^{N_1} \cdot \left(\frac{\Delta^\Delta}{(\Delta-1)^{\Delta-1}}\right)^{N_2}$$

such sequences.

Since we are counting decorated landscapes, the choice of $\textsc{Prev}(x)$ for each $x \in V(\mathcal{L})$ contributes the factor β^{N_2}.

Finally we need to specify the missing elements of the decoration, i.e. an element $\textsc{Final}(\mathcal{L}) \in b^{V(G_\mathcal{L})}$ and $\pi \in p^{V(G_\mathcal{L})}$ which contribute the factors b^{N_1} and p^{N_1} respectively. □

3.d. **The b-valued sequences encoded by a decorated landscape.**

Let $\mathcal{L}$ be a decorated landscape of height k. Let us inductively define *assignment functions* $\textsc{Asgn}_k, \textsc{Asgn}_{k-1}, \dots, \textsc{Asgn}_0 \colon V(G_\mathcal{L}) \to b$ as follows. First we let $\textsc{Asgn}_k = \textsc{Final}(\mathcal{L})$. Now suppose that $\textsc{Asgn}_i$ is defined for some $i \in \{1, \dots, k\}$ and let us define $\textsc{Asgn}_{i-1}$.

Let $y \in V(G_\mathcal{L})$. If there exists no $x \in V(\mathcal{L})_{i-1}$ such that $\overline{y} \in \mathrm{Var}(\overline{x})$ then we let $\textsc{Asgn}_{i-1}(y) := \textsc{Asgn}_i(y)$. Otherwise, since $\mathcal{L}$ is a landscape, there is a unique such x, and we let $\textsc{Asgn}_{i-1}(y) := \textsc{Prev}(x)(\overline{y})$.

Now for each $x \in V(G_\mathcal{L})$ we define a b-valued sequence $\textsc{Seq}(x)$ as follows. First let $\textsc{Seq}_0(x)$ be the empty sequence and suppose that $\textsc{Seq}_i(x)$ is defined for some $i < k$. If for some $y \in V(\mathcal{L})_i$ we have $\overline{x} \in \mathrm{Var}(\overline{y})$ then we define $\textsc{Seq}_{i+1}(x)$ to be the concatenation of $\textsc{Seq}_i(x)$ and the one element sequence consisting of $\textsc{Asgn}_{i+1}(x)$. Otherwise we let $\textsc{Seq}_{i+1}(x) := \textsc{Seq}_i(x)$. Finally, we let $\textsc{Seq}(x) := \textsc{Seq}_k(x)$.

To stress the dependence on $\mathcal{L}$ we might write $\textsc{Asgn}_i(\mathcal{L}, x)$ and $\textsc{Seq}(\mathcal{L}, x)$ instead of $\textsc{Asgn}_i(x)$ and $\textsc{Seq}(x)$, respectively. The following lemma is easy to check.

Lemma 3.8. *Let $\mathcal{K}$ and $\mathcal{L}$ be equivalent decorated landscapes. For all $x \in V(G_\mathcal{L}) = V(G_\mathcal{K})$ we have $\textsc{Seq}(\mathcal{K}, x) = \textsc{Seq}(\mathcal{L}, x)$.* □

3.e. **Landscape restrictions.**

Let $\mathcal{L}$ be a landscape and let H be a subgraph of $G_\mathcal{L}$. We consider H to be a variable graph by inducing the order on $N_H(x)$ and $N_G^{\mathrm{in}}(x)$ from $N_{G_\mathcal{L}}(x)$ and $N_{G_\mathcal{L}}^{\mathrm{in}}(x)$ respectively.

We proceed to define a new landscape $\mathrm{Res}_\mathcal{L}(H)$ referred to as the *restriction of $\mathcal{L}$ to H*. For brevity let us denote $\mathrm{Res}_\mathcal{L}(H)$ by $\mathcal{K}$. We set $G_\mathcal{K} := H$, and the local rule $\mathbf{R}_\mathcal{K}$ is defined as follows. If for $x \in V(H)$ we have $\mathrm{Var}_H(x) = \mathrm{Var}_{G_\mathcal{L}}(x)$ then we set $\mathbf{R}_\mathcal{K}(x) := \mathbf{R}_\mathcal{L}(x)$. Otherwise we let $\mathbf{R}_\mathcal{K}(x) := \mathcal{P}(b^{\mathrm{Var}_H(x)})$.

Finally we define $\mathrm{T}(\mathcal{K})$ as follows: for every tree τ in $T(\mathcal{L})$ we add to $\mathrm{T}(\mathcal{K})$ the connected components of the forest $\tau \cap H$. This finishes the definition of the pseudo-landscape $\mathcal{K}$. By construction it is a landscape. If $\mathcal{L}$ is decorated then we induce a decoration on $\mathcal{K}$ from $\mathcal{L}$ in the only sensible way.

We will say that the restriction $\mathrm{Res}_{\mathcal{L}}(H)$ is *faithful at a vertex* $x \in V(H)$ if $\mathrm{Var}_H(y) = \mathrm{Var}_{G_{\mathcal{L}}}(y)$ for every y such that $x \in \mathrm{Var}_{G_{\mathcal{L}}}(y)$. The following lemma is easy to check.

Lemma 3.9. *Let $\mathcal{K}$ be a restriction of $\mathcal{L}$ which is faithful at some vertex x. Then* $\textsc{Seq}_{\mathcal{K}}(x) = \textsc{Seq}_{\mathcal{L}}(x)$. □

If $V \subset V(G_{\mathcal{L}})$ then $\mathrm{Res}_{\mathcal{L}}(V)$ denotes the restriction to the graph induced on V from $G_{\mathcal{L}}$.

4. Local Lemma on Borel graphs of subexponential growth

4.a. **Analysis of the Moser-Tardos algorithm.**

Let $\mathcal{G}$ be a Borel graph and let $d\colon V(\mathcal{G}) \times V(\mathcal{G}) \to \mathbb{N} \cup \{\infty\}$ be the metric induced by the graph distance in $\mathcal{G}$. For $F \subset V(\mathcal{G})$ and $i \in \mathbb{N}$ we define F_{-i} to be the set of those $x \in F$ such that for any $y \notin F$ we have $d(x, y) \geqslant i$. Furthermore, for $x \in V(\mathcal{G})$ and $r \in \mathbb{N}$ we let $B(x, r) := \{y \in V(\mathcal{G})\colon d(x, y) \leqslant r\}$ be the ball of radius r around x. We say that a partition π of $V(\mathcal{G})$ is *r-sparse* if for every $x \in V(\mathcal{G})$ the different points of $B(x, r)$ belong to different parts of π.

The great majority of the work needed to obtain a Borel version of Local Lemma is contained in the following theorem. Recall that $X(\mathcal{G}, \mathbf{R})$ is defined to be the set $\{x \in V(\mathcal{G})\colon \mathbf{R}(x) \neq b^{\mathrm{Var}_{\mathcal{G}}(x)}\}$, i.e. it is the subset of $V(\mathcal{G})$ where the rule $\mathbf{R}$ is non-trivial.

Theorem 4.1. *Let $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, f)$ be a long Moser-Tardos tuple such that $\sup_{x \in V(\mathcal{G})} N_{\mathcal{G}}(x) < \infty$. Let us denote $D := \sup_{x \in X(\mathcal{G}, \mathbf{R})} |\mathrm{Var}(x)|$, $\Delta := \sup_{x \in V(\mathcal{G})} |N_{\mathrm{Rel}(\mathcal{G})}(x)|$, and $\beta := \sup_{x \in \Omega} |\mathbf{R}^c(x)|$, and let us assume that there are $\varepsilon \in (0, 1)$ and $n \in \mathbb{N}$ be such that the following conditions hold.*

1. *For any bounded Borel function $g\colon V(\mathcal{G}) \to \mathbb{N}$ supported on $X(\mathcal{G}, \mathbf{R})$ there exists a finite set $F \subset V(\mathcal{G})$ which is contained in $B(y, n)$ for some $y \in V(\mathcal{G})$, such that*

$$\max_{x \in V(\mathcal{G})} g(x) \;\leqslant\; \sum_{x \in F} g(x) \;<\; (1 + \varepsilon) \cdot \sum_{x \in F_{-3}} g(x). \tag{2}$$

2. *We have $\Delta \geqslant 2$ and*

$$b^{(1-\varepsilon) \cdot D} > \frac{\Delta^{\Delta}}{(\Delta - 1)^{\Delta - 1}} \cdot \beta, \tag{3}$$

3. *The partition π is n-sparse.*

Then there exist constants $K > 0$ and $M \in \mathbb{N}$ such that for all $k \in \mathbb{N}$ and all $N > M$, the probability that the k-step Moser-Tardos algorithm resamples some point more than N times is bounded by b^{-NK}.

Remark 4.2. (i) As explained in Remark 2.3, the implied probability space in the theorem above is the space $b^{p \times k}$ with the uniform probability measure.
(ii) We also note that $\sup_{x \in V(\mathcal{G})} N_{\mathcal{G}}(x) < \infty$ implies $\beta, D, \Delta < \infty$.

Proof of Theorem 4.1. We start by proving the following reduction.

Claim 4.3. *Without any loss of generality we may assume that for all $x \in X(\mathcal{G}, \mathbf{R})$ we have $|\operatorname{Var}(x)| = D$.*

Proof of Claim. If this is not the case then we construct a new long Moser-Tardos tuple $(\overline{\mathcal{G}}, \overline{\mathbf{R}}, \overline{\mathrm{I}}, \overline{\pi}, \overline{f})$ as follows.

We define $V(\overline{\mathcal{G}}) := V(\mathcal{G}) \cup (V(\mathcal{G}) \times \Delta)$, and

$$E(\overline{\mathcal{G}}) := E(\mathcal{G}) \cup \{(x, (x, i)) \colon x \in V(\mathcal{G}), i \geqslant |\operatorname{Var}_{\mathcal{G}}(x)|\}.$$

It is easy to see that $\overline{\mathcal{G}}$ is a Borel graph. It is also easy to see that $V(\mathcal{G})$ and $V(\mathcal{G}) \times \Delta$ are disjoint sets.[1] It follows that $\mathcal{G}$ is an induced subgraph of $\overline{\mathcal{G}}$, and by construction we see that for all $x \in V(\mathcal{G})$ we have $|\operatorname{Var}_{\overline{\mathcal{G}}}(x)| = D$.

Let $\overline{\mathbf{R}}$ be the Borel local rule for $\mathcal{G}$ defined as follows. For $x \in V(\mathcal{G})$ we let $f \in \overline{\mathbf{R}}(x)$ if and only if the restriction of f to $\operatorname{Var}_{\mathcal{G}}(x)$ is an element of $\mathbf{R}(x)$. For $x \notin V(\mathcal{G})$ we let $\overline{\mathbf{R}}(x)$ be the unique empty function (note that for $x \notin V(\mathcal{G})$ the set $\operatorname{Var}_{\overline{\mathcal{G}}}(x)$ is empty).

Note that $\operatorname{Rel}(\mathcal{G})$ is an induced subgraph of $\operatorname{Rel}(\overline{\mathcal{G}})$, and that its complement is a Borel graph with no edges. This allows to easily extend the independence function I for $\operatorname{Rel}(\mathcal{G})$ to an independence function $\overline{\mathrm{I}}$ for $\operatorname{Rel}(\overline{\mathcal{G}})$.

Recall that p denotes the number of parts in the partition π. We extend the partition π of $V(\mathcal{G})$ to a partition $\overline{\pi}$ of $V(\overline{\mathcal{G}})$ by adding to it sets $W_j \times \{i\}$, for $j \in p$ and $i \in \Delta$, in an arbitrary order.

It is straightforward to see that the tuple $(\overline{\mathcal{G}}, \overline{\mathbf{R}}, \overline{\mathrm{I}}, \overline{\pi}, \overline{f})$ fulfils all the conditions of Theorem 4.1, and additionally for all $x \in X(\mathcal{G}, \mathbf{R})$ we have $|\operatorname{Var}_{\overline{\mathcal{G}}}(x)| = D$. Now the statement of the claim follows from Lemma 2.4. □

Thus from now on we assume that for all $x \in X(\mathcal{G}, \mathbf{R})$ we have $|\operatorname{Var}(x)| = D$. Let us define, for all $k \in \mathbb{N}$ and all $\mathtt{rnd} \in b^{p \times k}$, a decorated landscape $\mathcal{L} = \mathcal{L}(\mathtt{rnd})$ using the functions $\mathrm{M}_0(\mathtt{rnd}), \ldots, \mathrm{M}_k(\mathtt{rnd})$ given by the k-step Moser-Tardos algorithm.

We let $G_{\mathcal{L}} := \mathcal{G}$, $\mathbf{R}_{\mathcal{L}} := \mathbf{R}$, and the forest in $\operatorname{Canvas}(G_{\mathcal{L}})$ is defined as follows. For $i = 0, \ldots, k-1$ we let $V(\mathcal{L})_i := \mathrm{I}\,\mathrm{B}(M_i) \times \{i\}$. We claim that for $i \in k-1$ and $x \in V(\mathcal{L})_{i+1}$ there exists $y \in V(\mathcal{L})_i$ such that $\overline{x}$ and $\overline{y}$ are connected in $\operatorname{Rel}(\mathcal{G})$. Indeed, otherwise it is easy to check that x must be in $B(M_i)$ which contradicts the maximality property of $IB(M_i)$.

Now given $x \in V(\mathcal{L})_{i+1}$ let z be the vertex in $V(\mathcal{L})_i$ such that $\overline{z} \leqslant \overline{y}$ for all $y \in V(\mathcal{L})_i$ for which $\overline{x}$ and $\overline{y}$ are connected in $\operatorname{Rel}(\mathcal{G})$. In particular there is an edge from z to x in $\operatorname{Canvas}(G_{\mathcal{L}})$ and we add this edge to our forest.

[1] To prove it one needs to fix a definition of an ordered pair. We use the most standard definition due to Kuratowski. An alternative would be to explicitly define $V(\overline{\mathcal{G}})$ to be a disjoint union of $V(\mathcal{G})$ and $V(\mathcal{G}) \times \Delta$.

The decoration is defined as follows. The function $\pi\colon V(\mathcal{G}) \to \mathbb{N}$ is already defined; we let $\textsc{Final}(\mathcal{L}) := \mathrm{M}_k$ and for $x = (\overline{x}, \ell(x)) \in V(\mathcal{L})$ we let $\textsc{Prev}(x)$ to be the function $\mathrm{M}_{\ell(x)}$ restricted to $\mathrm{Var}(x)$.

The following claim is easy to check.

Claim 4.4. *For $x \in V(\mathcal{G})$ we have $\textsc{Seq}_{\mathcal{L}(\mathtt{rnd})}(x) = \textsc{Used}(\mathtt{rnd}, x)$.* □

Let $g(\mathtt{rnd})\colon V(\mathcal{G}) \to b$ be the function $x \mapsto |V(\mathcal{L}) \cap (\{x\}\times\mathbb{N})|$. It is routine to check that $g(\mathtt{rnd})$ is Borel. Furthermore, clearly for $\mathtt{rnd} \in b^{p\times k}$ we have that $g(\mathtt{rnd})$ is bounded by k. Let $F(\mathtt{rnd}) \subset V(\mathcal{G})$ be the set guaranteed for the function $g(\mathtt{rnd})$ by the first assumption in the statement of Theorem 4.1.

Recall that $\mathbb{S}$ is the set of all finite b-valued sequences. Let L be the set of (the iso-classes of) the grounded landscapes, including the empty landscape. The length of $s \in \mathbb{S}$ will be denoted by $|s|$. If $h\colon p \to \mathbb{S}$ then we let $\mathrm{concat}(h)$ be the concatenation of the sequences $h(0), h(1), \ldots, h(p-1)$.

We proceed to define an injective map

$$\zeta\colon b^{p\times k} \to \mathcal{P}(p) \times \mathbb{S} \times L.$$

If $|V(\mathcal{L}(\mathtt{rnd}))|$ is empty then we let

$$\zeta(\mathtt{rnd}) := (\emptyset, \mathrm{concat}(\mathtt{rnd}), \text{empty landscape}).$$

Otherwise we let $\mathcal{K}(\mathtt{rnd})$ be a grounded landscape equivalent to $\mathrm{Res}_{\mathcal{L}(\mathtt{rnd})}(F(\mathtt{rnd}))$, and we define

$$\zeta(\mathtt{rnd}) := (\pi(F(\mathtt{rnd})_{-2}), \mathrm{concat}(U(\mathtt{rnd})), \mathcal{K}(\mathtt{rnd})),$$

where $U(\mathtt{rnd})$ is defined as follows: if $i \in \pi(F(\mathtt{rnd})_{-2})$ then we let $U(\mathtt{rnd})(i)$ be the sequence $\textsc{Unused}(\mathtt{rnd}, x)$, and otherwise we let $U(\mathtt{rnd})(i)$ be the sequence $\mathtt{rnd}(i)$.

Let us argue that ζ is indeed an injection. First note that for any set F the restriction of $\mathcal{L}(\mathtt{rnd})$ to F is faithful at all vertices in F_{-2}. By Lemmas 3.8 and 3.9 It follows that if $y \in F(\mathtt{rnd})_{-2}$ then $\textsc{Seq}_{\mathcal{K}(\mathtt{rnd})}(y)$ is equal to $\textsc{Used}(\mathtt{rnd}, y)$.

Let us argue that we can recover the function $U(\mathtt{rnd})$ from the triple $(\pi(F(\mathtt{rnd})_{-2}), \mathrm{concat}(U(\mathtt{rnd})), \mathcal{K}(\mathtt{rnd}))$. It is enough to show that we can recover the lengths $|U(\mathtt{rnd})(i)|$ for $i \in p$ from the pair $(\pi(F(\mathtt{rnd})_{-2}), \mathcal{K}(\mathtt{rnd}))$. This is possible since since π assigns different values to all the vertices of $G_{\mathcal{K}(\mathtt{rnd})}$. In particular knowing $\pi(F(\mathtt{rnd})_{-2})$ allows us to deduce which vertices of $G_{\mathcal{K}(\mathtt{rnd})}$ belong to $F(\mathtt{rnd})_{-2}$, and so we can recover the lengths $|U(\mathtt{rnd})(i)|$ as follows:

$$|U(\mathtt{rnd})(i)| = \begin{cases} k - |\textsc{Seq}_{\mathcal{K}(\mathtt{rnd})}(y)| & \text{if } i \in \pi(F(\mathtt{rnd})_{-2}), \\ k & \text{otherwise} \end{cases}$$

Now it is easy to recover $\mathtt{rnd}$. For $i \notin \pi(F(\mathtt{rnd})_{-2})$ we have $\mathtt{rnd}(i) = U(\mathtt{rnd})(i)$, and for $x \in F(\mathtt{rnd})_{-2}$ we have that $\mathtt{rnd}(\pi(y))$ is a concatenation of $\textsc{Unused}(\mathtt{rnd}, y) = U(\mathtt{rnd})(\pi(y))$ and $\textsc{Used}(\mathtt{rnd}, y) = \textsc{Seq}_{\mathcal{K}(\mathtt{rnd})}(y)$. This establishes that the map ζ is an injection.

Let us denote $N_1 := \max_{y\in V(\mathcal{G})} |B(y,n)|$. Since we assume $\sup_{x\in V(\mathcal{G})} N_{\mathcal{G}}(x) < \infty$, we have that N_1 is finite.

Claim 4.5. *Let $N_2 \in \mathbb{N}$. Let $\mathcal{K}$ be a grounded landscape of type $(D,\Delta,\beta,N_1,N_2,p)$ and let $A \in \mathcal{P}(p)$. The number of points in $\mathrm{im}(\zeta)$ with the first and the third coordinate equal to, respectively, A and $\mathcal{K}$ is bounded by*

$$b^{pk-(1-\varepsilon)N_2\cdot D}.$$

Proof of Claim. We are counting the possible sequences $\mathrm{concat}(U(\mathtt{rnd}))$ for those $\mathtt{rnd} \in b^{p\times k}$ such that $\pi(F(\mathtt{rnd})_{-2}) = A$ and $\mathcal{K}(\mathtt{rnd}) = \mathcal{K}$.

Clearly $\mathrm{concat}(U(\mathtt{rnd}))$ is a b-valued sequence of length equal to $p\cdot k$ minus "the number of bits used by the k-step Moser-Tardos algorithm for resampling points in $F(\mathtt{rnd})_{-2}$", i.e.

$$\sum_{x\in F(\mathtt{rnd})_{-2}} h_k(x).$$

The latter sum is easily seen to be at least $|V(\mathcal{L}) \cap (F(\mathtt{rnd})_{-3}\times\mathbb{N})|$ times D. This is because whenever we have a point $y \in V(\mathcal{L}) \cap (F(\mathtt{rnd})_{-3}\times\mathbb{N})$, we have that in the Moser-Tardos algorithm's passage from $\mathrm{M}_{\ell(y)}$ to $\mathrm{M}_{\ell(y)+1}$ all the elements of $\mathrm{Var}(\overline{y})$ got resampled, and clearly $\mathrm{Var}(\overline{y}) \subset F(\mathtt{rnd})_{-2}$.

The number of points in $V(\mathcal{L}) \cap (F(\mathtt{rnd})_{-3}\times\mathbb{N})$ is, by the inequality (4), at least $\frac{1}{1+\varepsilon}N_2 \geqslant (1-\varepsilon)N_2$. Thus all in all we see that $\mathrm{concat}(U(\mathtt{rnd}))$ is a b-valued sequence of length at most $pk-(1-\varepsilon)N_2\cdot D$, which establishes the lemma. □

For $N_2 \in \mathbb{N}$ let $P(N_2)$ be the probability that $\mathtt{rnd} \in b^{p\times k}$ is such that $V(\mathcal{K}(\mathtt{rnd})) = N_2$. Since ζ is an injection we have that $P(N_2)\cdot b^{p\times k}$ is equal to the number of elements in $\mathrm{im}(\zeta)$ whose fourth coordinate is a grounded landscape of type $(D,\Delta,\beta,N_1,N_2,p)$. This, together with Lemma 3.7 and Claim 4.5 allows us to bound $P(N_2)$ from above. After taking base-b logarithms, denoted by log, we obtain

$$\begin{aligned}\log(P(N_2)\cdot b^{pk}) = \log(P(N_2)) + pk \leqslant \\ \leqslant pk - (1-\varepsilon)N_2\cdot D + \log(C) + N_1\log(N_2) + N_2\log(\frac{\Delta^\Delta}{(\Delta-1)^{\Delta-1}}\cdot\beta)\end{aligned}$$

where C is the constant from Lemma 3.7. In particular, C does not depend on N_2.

After subtracting pk from both sides and dividing both sides by N_2, we obtain

$$\frac{1}{N_2}\log(P(N_2)) \leqslant \frac{\log(C)+N_1\log(N_2)}{N_2} - (1-\varepsilon)D + \log(\frac{\Delta^\Delta}{(\Delta-1)^{\Delta-1}}\cdot\beta),$$

Note that the term $\frac{\log(C)+N_1\log(N_2)}{N_2}$ converges to 0 as $N_2\to\infty$. Because of that and the assumption (6), we see that there exist constants $\overline{K} > 0$ and $M \in \mathbb{N}$, both independent of N_2, such that for all $N_2 > M$ we have

$$\log(P(N_2)) \leqslant -N_2\cdot\overline{K},$$

and so $P(N_2) \leqslant b^{-N_2 \cdot \overline{K}}$.

Let $Q(N_2)$ be the probability that $\mathtt{rnd} \in b^{p \times k}$ is such that $|V(\mathcal{K}(\mathtt{rnd}))| > N_2$. Clearly $Q(N_2) = \sum_{n_2 > N_2} P(n_2)$, and after making $\overline{K}$ smaller if necessary we see that for $N_2 > M$ we have $Q(N_2) \leqslant b^{-N_2 \cdot \overline{K}}$.

The claim of the theorem now easily follows: Let $G \subset b^{p \times k}$ be the set of those $\mathtt{rnd}$ such that $V(\mathcal{K}(\mathtt{rnd})) > N_2$. We just showed that the measure of G is at most $b^{-N_2 \cdot \overline{K}}$. We claim that if $\mathtt{rnd} \notin G$ then the k-step Moser-Tardos algorithm resamples each point less than $\Delta \cdot N_2$ times.

Indeed, let us assume that this is not the case and let $x \in V(\mathcal{G})$ be a point which is resampled at least $\Delta \cdot N_2$ times. However, since $\Delta = \sup_{x \in V(\mathcal{G})} |N_{\mathrm{Rel}(\mathcal{G})}(x)|$, we see that in particular there are at most Δ points $y \in V(\mathcal{G})$ such that $x \in \mathrm{Var}(y)$. It follows that for one of these points y we have $|V(\mathcal{L}(\mathtt{rnd})) \cap \{y\} \times \mathbb{N}| \geqslant N_2$. Hence, by by the choice of the set $F(\mathtt{rnd})$, we have $|V(\mathcal{K}(\mathtt{rnd}))| \geqslant N_2$, which is a contradiction with the fact that $\mathtt{rnd} \notin G$.

Thus in order to finish the proof of Theorem 4.1 we would let $K := \frac{\overline{K}}{\Delta}$ if we only cared about numbers N which are multiples of Δ. It is straightforward to see that taking $K := \frac{\overline{K}}{2\Delta}$ works for all N. □

4.b. **Borel Local Lemma.**

Let us start by showing that the first item of Theorem 4.1 is automatically fulfilled for arbitrary ε and large enough n if $\mathcal{G}$ is of uniformly subexponential growth.

Lemma 4.6. *Let $\varepsilon > 0$ and let us assume that $n \in \mathbb{N}$ is such that for all $x \in V(\mathcal{G})$ we have $B(x, 3n) < (1+\varepsilon)^n$. Let $g \colon V(\mathcal{G}) \to \mathbb{N}$ be a bounded Borel function which assumes its maximum at some point $y \in V(\mathcal{G})$. Then there exists $r \in \{3, \dots, 3n\}$ such that*

$$\sum_{x \in B(y,r)} g(x) < (1+\varepsilon) \cdot \sum_{x \in B(y,r-3)} g(x). \tag{4}$$

Proof. Suppose that for all $r \in \{3, \dots, 3n\}$ the inequality (4) does not hold. Then

$$\sum_{x \in B(y,3n)} g(x) \geqslant (1+\varepsilon)^n \cdot g(y).$$

Since $B(y, 3n) < (1+\varepsilon)^n$, we deduce that for some $z \in B(x, 3n)$ we have $g(z) > g(y)$, which contradicts the maximality of $g(y)$. □

We are ready to give a criterion when the functions returned by the Moser-Tardos algorithm converge pointwise to a Borel function.

Theorem 4.7. *Let $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, f)$ be a long Moser-Tardos tuple such that $\mathcal{G}$ has a uniformly subexponential growth. Let Δ be the maximal degree in $\mathrm{Rel}(\mathcal{G})$, and let us assume*

that $\Delta \geqslant 2$ and that for all $x \in V(\mathcal{G})$ we have

$$1 - \frac{|\mathbf{R}(x)|}{|b^{\mathrm{Var}(x)}|} < \frac{(\Delta-1)^{\Delta-1}}{\Delta^{\Delta}}. \tag{5}$$

If π is m-sparse for sufficiently large m then there exists a constant $K > 0$ such that the probability that the $\mathbb{N}$-step Moser-Tardos algorithm for $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, f)$ resamples some point more than N times is bounded by b^{-NK}.

In particular, the functions $\mathrm{M}_0, \mathrm{M}_1, \ldots$ returned by the Moser-Tardos algorithm converge with probability 1 to a Borel function which satisfies $\mathbf{R}$.

Proof. Let us first explain the "In particular" part. Note that the functions M_i, $i \in \mathbb{N}$ are bounded Borel functions. Thus if their pointwise limit exists, it is automatically a Borel function.

By the main part of the theorem we have that with probability 1 there exists a constant N such that the $\mathbb{N}$-step Moser-Tardos Algortihm resamples at most N times. Since $\mathcal{G}$ is of uniformly subexponential growth, in particular balls of finite radii in $V(\mathcal{G})$ are finite.

Thus for every $x \in V(\mathcal{G})$ there exists i such that for all $j \geqslant i$ we have that $B(x,4) \cap \mathrm{I}\,\mathrm{B}(\mathrm{M}_j) = \emptyset$. This shows that in fact $B(x,2) \cap \mathrm{B}(\mathrm{M}_j) = \emptyset$ for all $j \geqslant i$. This implies that for $j \geqslant i$ we have that the restriction of M_j and M_i to $B(x,2)$ are equal. This together with the fact that $x \notin \mathrm{B}(\mathrm{M}_j)$ for $j \geqslant i$ establishes the "In particular" part.

Claim 1. *In order to prove the theorem it is enough to establish the existence of a constant $K > 0$ with the property that for every k the probability that there exists $x \in V(\mathcal{G})$ such that the k-step Moser-Tardos algorithm resamples x more than N times is bounded by b^{-NK}.*

Proof of Claim. For $i \in \mathbb{N}$ the product measure on on $b^{p \times i}$ will be denoted by μ_i, and the product measure on $b^{p \times \mathbb{N}}$ will be denoted by μ. Let $\pi_i \colon b^{p \times \mathbb{N}} \to b^{p \times i}$ be the restriction to the first i coordinates.

Let $U_i \subset b^{p \times i}$ be the set of those `rnd` for which the i-step Moser-Tardos algorithm resamples each point at most N times. Let us assume that for all $i \in \mathbb{N}$ we have $\mu_b(U_i) \geqslant 1 - b^{-NK}$.

Let $V_i := \pi_a^{-1}(U_i)$. Note that we have $\mu(V_i) \geqslant 1 - b^{NK}$ and for $j > i$ we have $V_j \subset V_i$. It follows that $\mu(\bigcap_{i \in \mathbb{N}} V_i) \geqslant 1 - b^{NK}$ and the claim follows since for any $\mathtt{rnd} \in \bigcap_{i \in \mathbb{N}} V_i$ the $\mathbb{N}$-step Moser-Tardos algorithm resamples each point at most N times. □

Thus we proceed to show the existence of a constant $K > 0$ such that for every k the probability that the k-step Moser-Tardos algorithm resamples some point more than N times is bounded by b^{-NK}.

Let $D := \sup_{x \in V(\mathcal{G})} |\,\mathrm{Var}(x)|$ and $\beta := \sup_{x \in \Omega} |\mathbf{R}^c(x)|$. Since we assume (5), we also have

$$\frac{\beta}{b^D} < \frac{(\Delta-1)^{\Delta-1}}{\Delta^{\Delta}}.$$

Thus we can fix a number $\varepsilon \in (0,1)$ such that

$$b^{(1-\varepsilon)\cdot D} > \frac{\Delta^\Delta}{(\Delta-1)^{\Delta-1}}\cdot\beta. \tag{6}$$

Since $\mathcal{G}$ is of uniformly subexponential growth, we can fix a number $n \in \mathbb{N}$ such that

$$|B(x,3n)| < (1+\varepsilon)^n.$$

Let us assume that π is $3n$-sparse. We claim that all the conditions of Theorem 4.1 are fulfilled with ε and $3n$. Indeed, the first item of Theorem 4.1 follows from Lemma 4.6 and the second one follows from the inequality (6). Thus we can apply Theorem 4.1, which finishes the proof of Theorem 4.7. □

Corollary 4.8 (Borel Local Lemma)**.** *Let $\mathcal{G}$ be a Borel variable graph and let $\mathbf{R}$ be a Borel local rule on $\mathcal{G}$. Furthermore let us assume that the graph $\mathrm{Rel}(\mathcal{G})$ is of uniformly subexponential growth, and let Δ be the maximal degree in $\mathrm{Rel}(\mathcal{G})$. If for all $x \in V(\mathcal{G})$ we have*

$$1 - \frac{|\mathbf{R}(x)|}{|b^{\mathrm{Var}(x)}|} < \frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}.$$

then there exists a Borel function $f\colon V(G) \to b$ which satisfies $\mathbf{R}$.

Proof. By Lemma A.3 in the appendix, there exists an independence function for $\mathrm{Rel}(\mathcal{G})$ so we let I be such an independence function. Let $m \in \mathbb{N}$ be as in the previous theorem. By Lemma A.4 om the appendix, we can let π to be an m-sparse partition of $V(\mathcal{G})$. Finally let $f\colon V(\mathcal{G}) \to b$ be a constant function equal to 0 everywhere.

If $\Delta = 1$ then it means that the sets $\mathrm{Var}(x)$, $x \in V(\mathcal{G})$, are pairwise disjoint. In this case it is very easy to see that the functions $\mathrm{M}_0, \mathrm{M}_1, \dots$ returned by the $\mathbb{N}$-step Moser-Tardos algorithm for $(\mathcal{G}, \mathbf{R}, \mathrm{I}, \pi, f)$ converge with probability 1 to a Borel function which satisfies $\mathbf{R}$. As always, probability is with respect to the choice of the random seed $\mathtt{rnd} \in b^{p\times\mathbb{N}}$

If $\Delta \geqslant 2$ then $\mathrm{M}_0, \mathrm{M}_1, \dots$ also converge with probability 1 to a Borel function which satisfies $\mathbf{R}$, by Theorem 4.7. □

Appendix A. Some standard lemmas

We provide this appendix only for reader's convenience, as all the results are standard. We start with the following lemma.

Lemma A.1. *Let Ω be a standard Borel space and let (U,γ) be a Borel arrow on Ω. There exists a partition of U into disjoint Borel sets $U_0, U_1, \dots$, such that for $x \in U_0$ we have $\gamma(x) = x$ and for all $i > 0$ and $x \in U_i$ we have $\gamma(x) \notin U_i$.*

Proof. By Kuratowski's theorem on standard Borel spaces we might as well assume that $\Omega \subset [0,1)$. Let $d\colon \Omega\times\Omega \to [0,1)$ be the standard metric on Ω. Let $U_0 := \{x \in U\colon \gamma(x) = x\}$, and let $V := U \setminus U_0$. Let $V_n := \{x \in V\colon d(x,\gamma(x)) > \frac{1}{n}\}$. Let us fix n and let $W := V_n$.

Clearly it is enough to partition W into countably many sets W_i, $i = 0, 1, \dots$ such that for $x \in W_i$ we have $\gamma(x) \notin W_i$.

For $j \in n$ let $A_j := [\frac{j}{n}, \frac{j+1}{n})$. Note that $\gamma(A_j) \cap A_j = \emptyset$. Thus we also have $\gamma(A_j \cap W) \cap A_j \cap W = \emptyset$, and we may define $W_j := A_j \cap W$. □

Corollary A.2. *Let $\mathcal{G}$ be a locally finite Borel graph without self-loops. There exists a sequence $S = ((V_0, \gamma_0), (V_1, \gamma_1), \dots)$ of Borel arrows on $V(\mathcal{G})$ such that $\mathcal{G} = \mathcal{G}(S)$, and with the property that for all $i \in \mathbb{N}$ we have $N_{\mathcal{G}}(U_i) \cap U_i = \emptyset$.*

Proof. The previous lemma, together with the assumption about the self loops, show that we can find a sequence $S = ((U_0, \gamma_0), (U_1, \gamma_1), \dots)$ of Borel arrows such that $\mathcal{G} = \mathcal{G}(S)$, and such that for all $i \in \mathbb{N}$ we have $U_i \cap \gamma_i(U_i) = \emptyset$.

Let $\mathbf{P}$ be the set of all finite sets of natural numbers. Let $M \colon V(\mathcal{G}) \to \mathbf{P}$ be defined as follows. For $x \in V(\mathcal{G})$ we let $M(x)$ be the set of those i such that $x \in U_i$ and there is no $j < i$ such that $x \in U_j$ and $\gamma_j(x) = \gamma_i(x)$. It is not difficult to check that M is a Borel function (i.e. for $s \in \mathbf{P}$ we have that $M^{-1}(s)$ is a Borel set).

Now for $i \in \mathbb{N}$ and $s \in \mathbf{P}$ we define $U(i, s) := U_i \cap M^{-1}(s)$ and let $\gamma(i, s)$ be the restriction of γ_i to $U(i, s)$. Clearly the sets $U(i, s)$ are Borel, and the family $\overline{S}$ of Borel arrows $((U(i, s), \gamma(i, s)))$, $i \in \mathbb{N}$, $s \in \mathbf{P}$ has the property that $\mathcal{G} = \mathcal{G}(S) = \mathcal{G}(\overline{S})$.

Let us argue that for $x \in U(i, s)$ we have $N_{\mathcal{G}}(x) \cap U(i, s) = \emptyset$. By the definition of M we have $U(i, s) \subset \bigcap_{j \in s} U_j$ and $N_{\mathcal{G}}(x) = \{\gamma_j(x) \colon j \in s\}$. Now the lemma follows because for all $i \in \mathbb{N}$ we have $U_i \cap \gamma_i(U_i) = \emptyset$. □

Let $\mathcal{G}$ be Borel graph and let $\mathcal{B}$ be the Borel σ-algebra of $V(\mathcal{G})$. Recall that an independence function for $\mathcal{G}$ is a function $\mathrm{I} \colon \mathcal{B} \to \mathcal{B}$ such that for $X \in \mathcal{B}$ the set $\mathrm{I}(X)$ is a maximal subset of X which is independent in $\mathcal{G}$.

Lemma A.3. *Let $\mathcal{G}$ be a locally finite Borel graph. There exists an independence function* I *for $\mathcal{G}$.*

Proof. Let us first note that that Lemma A.1 implies in particular that the operation of removing all self-loops from $\mathcal{G}$ results in a *Borel* graph. Note also that if I is an independence function for "$\mathcal{G}$ with self-loops removed" then it is also an independence function for $\mathcal{G}$. Thus we may just as well assume that $\mathcal{G}$ does not have self-loops.

Furthermore, we have that I is an independence function for $\mathcal{G}$ if and only if it is an independence function for the symmetrization of $\mathcal{G}$, so we may assume that $\mathcal{G}$ is symmetric.

If $V(\mathcal{G})$ is countable we can use the following greedy algorithm: first we fix an enumeration $x_0, x_1, \dots$ of $V(\mathcal{G})$ and for $X \subset V(\mathcal{G})$ we define $\mathrm{I}(X)$ inductively as follows. We let X_0 be equal to $\{x_0\} \cap X$. Assume that X_i is defined for some $i \in \mathbb{N}$ and let $X_{i+1} := X_i \cup ((\{x_{i+1}\} \cap X) \setminus N(X_i))$. Finally we let $\mathrm{I}(X) := \bigcup_{i \in \mathbb{N}} X_i$. It is clear that I has all the required properties.

Corollary A.2 allows us to adapt this algorithm to an arbitrary Borel graph $\mathcal{G}$. Indeed, let $S = ((V_0, \gamma_0), (V_1, \gamma_1), \dots)$ be a sequence of Borel arrows as in Corollary A.2, and let $X \in \mathcal{B}$. We let $X_0 := V_0 \cap X$ and we assume that X_i is defined for some $i \in \mathbb{N}$. Now we define $X_{i+1} := X_i \cup ((V_{i+1} \cap X) \setminus N(X_i))$ and $\mathrm{I}(X) := \bigcup_{i\in\mathbb{N}} X_i$.

It is clear that $\mathrm{I}(X)$ is a Borel subset of X. The independence of $\mathrm{I}(X)$ follows from Corollary A.2. Indeed it is enough to inductively show that the sets $X_0, X_1, \dots$ are independent sets. For X_0 it follows directly from Corollary A.2, so let assume that we we know it for some i.

Consider the graph on the vertex set X_{i+1} induced from $\mathcal{G}$. We need to show that it has no edges. By definition we have $X_{i+1} := X_i \cup ((V_{i+1} \cap X) \setminus N(X_i))$. There are no edges between points of X_i and X_i by the inductive assumption. There are also no edges between points of $(V_{i+1} \cap X) \setminus N(X_i)$ and $(V_{i+1} \cap X) \setminus N(X_i)$ by Corollary A.2. Finally there are clearly no edges between X_i and $(V_{i+1} \cap X) \setminus N(X_i)$. This finishes the proof of the independence of $\mathrm{I}(X)$.

Let us finally argue about the maximality of $\mathrm{I}(X)$. Suppose there is $x \in X \setminus \mathrm{I}(X)$ such that $\mathrm{I}(X) \cup \{x\}$ is independent. Let $i \in \mathbb{N}$ be the smallest natural number such that $x \in V_i$. If $i = 0$ then clearly $x \in X_0$. Thus let us assume $i > 0$. Since $\mathrm{I}(X) \cup \{x\}$ is independent we have $x \notin N(X_{i-1})$ and so $x \in X_i$. Since $X_i \subset \mathrm{I}(X)$ we have $x \in \mathrm{I}(X)$ which finishes the proof. □

Lemma A.4. *Let $\mathcal{G}$ be a Borel graph such that $\sup_{x\in V(\mathcal{G})} |N(x)| < \infty$ and let $r \in \mathbb{N}$. There exists a Borel partition of $V(\mathcal{G})$ which is r-sparse.*

Proof. Let $\mathcal{H}$ be the Borel graph defined as follows. We let $V(\mathcal{H}) := V(\mathcal{G})$ and we let $(x, y) \in E(\mathcal{H})$ if the distance between x and y in $\mathcal{G}$ is at most n. Let I be an independence function for $\mathcal{H}$.

We proceed to inductively define a sequence of pair-wise disjoint Borel subsets $W_0, W_1, \dots$ of $V(\mathcal{H})$ as follows: let $W_0 := \mathrm{I}(V(\mathcal{G}))$, and let us assume that for some $i \in \mathbb{N}$ all the sets $W_0, \dots, W_{i-1}$ are defined. Then we let

$$W_i := \mathrm{I}(V(\mathcal{G}) \setminus \bigcup_{j=0}^{i-1} W_j).$$

Clearly it is enough to argue that for some n we have $\bigcup_{j=0}^{n-1} W_j = V(\mathcal{G})$. Let $m := \sup_{x\in V(\mathcal{G})} |B_{\mathcal{H}}(x, 2)|$. By the assumption on $\mathcal{G}$ we have $m < \infty$, and we claim that $\bigcup_{j=0}^{m-1} W_j = V(\mathcal{G})$.

To show it, it is enough to argue that for every $x \in V(\mathcal{H})$ and every $j \in m$ we have that either (i) $W_j \cap B_{\mathcal{H}}(x, 2) \neq \emptyset$, or (ii) $B_{\mathcal{H}}(x, 1) \subset \bigcup_{i=0}^{j-1} W_i$.

If for some j the condition (ii) does not hold then clearly

$$V(\mathcal{G}) \setminus \bigcup_{j=0}^{i-1} W_j \quad \cap \quad B_{\mathcal{H}}(x, 1) \neq \emptyset.$$

But $\mathrm{I}(V(\mathcal{G}) \setminus \bigcup_{j=0}^{i-1} W_j)$ is a maximal independent set of $V(\mathcal{G}) \setminus \bigcup_{j=0}^{i-1} W_j$ and hence must contain a point in $B_{\mathcal{H}}(x, 2)$, as claimed. This finishes the proof.

□

References

[1] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., Hoboken, NJ, third edition, 2008. With an appendix on the life and work of Paul Erdős.

[2] A. Bernshteyn. Measurable versions of the lovász local lemma and measurable graph colorings. Apr 2016. Available at http://arxiv.org/abs/1604.07349v2.

[3] P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday), Vol. II*, pages 609–627. Colloq. Math. Soc. János Bolyai, Vol. 10. North-Holland, Amsterdam, 1975.

[4] G. Kun. Expanders have a spanning lipschitz subgraph with large girth. Mar 2013. Available at http://arxiv.org/abs/1303.4982v2.

[5] R. A. Moser. A constructive proof of the Lovász local lemma. In *STOC'09—Proceedings of the 2009 ACM International Symposium on Theory of Computing*, pages 343–350. ACM, New York, 2009.

[6] R. A. Moser and G. Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2):Art. 11, 15, 2010.

[7] J. B. Shearer. On a problem of Spencer. *Combinatorica*, 5(3):241–245, 1985.

[8] J. H. Spencer. Robin moser makes Lovász Local Lemma Algorithmic! Notes of Joel Spencer. Available at https://www.cs.nyu.edu/spencer/moserlovasz1.pdf.

[9] J. H. Spencer. Asymptotic lower bounds for Ramsey functions. *Discrete Math.*, 20(1):69–76, 1977/78.